arXiv:1405.4917v2 [cs.LO] 2 Jun 2014

An Algebraic Hardness Criterion for Surjective Constraint Satisfaction

Hubie Chen
Departamento LSI
Facultad de Informática
Universidad del País Vasco
San Sebastián, Spain
and
IKERBASQUE, Basque Foundation for Science
Bilbao, Spain

Abstract

The constraint satisfaction problem (CSP) on a relational structure $\mathbf{B}$ is to decide, given a set of constraints on variables where the relations come from $\mathbf{B}$, whether or not there is a assignment to the variables satisfying all of the constraints; the surjective CSP is the variant where one decides the existence of a surjective satisfying assignment onto the universe of $\mathbf{B}$. We present an algebraic condition on the polymorphism clone of $\mathbf{B}$ and prove that it is sufficient for the hardness of the surjective CSP on a finite structure $\mathbf{B}$, in the sense that this problem admits a reduction from a certain fixed-structure CSP. To our knowledge, this is the first result that allows one to use algebraic information from a relational structure $\mathbf{B}$ to infer information on the complexity hardness of surjective constraint satisfaction on $\mathbf{B}$. A corollary of our result is that, on any finite non-trivial structure having only essentially unary polymorphisms, surjective constraint satisfaction is NP-complete.

1 Introduction

The *constraint satisfaction problem (CSP)* is a computational problem in which one is to decide, given a set of constraints on variables, whether or not there is an assignment to the variables satisfying all of the constraints. This problem appears in many guises throughout computer science, for instance, in database theory, artificial intelligence, and the study of graph homomorphisms. One obtains a rich and natural family of problems by defining, for each relational structure $\mathbf{B}$, the problem $\mathsf{CSP}(\mathbf{B})$ to be the case of the CSP where the relations used to specify constraints must come from $\mathbf{B}$. An increasing literature studies the algorithmic and complexity behavior of this problem family, focusing on finite and finite-like structures [1, 12, 2]; a primary research issue is to determine which such problems are polynomial-time tractable, and which are not. To this end of classifying problems, a so-called *algebraic approach* has been quite fruitful [5]. In short, this approach is founded on the facts that the complexity of a problem $\mathsf{CSP}(\mathbf{B})$ depends (up to polynomial-time reducibility) only on the set of relations that are *primitive positive definable* from $\mathbf{B}$, and that this set of relations can be derived from the *clone of polymorphisms* of $\mathbf{B}$. Hence, the project of classifying all relational structures according to the complexity of $\mathsf{CSP}(\mathbf{B})$ can be formulated as a classification question on clones;

this permits the employment of algebraic notions and techniques in this project. (See the next section for formal definitions of the notions discussed in this introduction.)

A natural variant of the CSP is the *surjective CSP*, where an instance is again a set of constraints, but one is to decide whether or not there is a *surjective* satisfying assignment to the variables. For each relational structure $\mathbf{B}$, one may define $\mathsf{SCSP}(\mathbf{B})$ to be the surjective CSP on $\mathbf{B}$, in analogy to the definition of $\mathsf{CSP}(\mathbf{B})$. Note that one can equivalently define $\mathsf{SCSP}(\mathbf{B})$ to be the problem of deciding, given as input a relational structure $\mathbf{A}$, whether or not there is a surjective homomorphism from $\mathbf{A}$ to $\mathbf{B}$. An early result on this problem family was the complexity classification of all two-element structures [7, Proposition 6.11], [8, Proposition 4.7]. There is recent interest in understanding the complexity of these problems, which perhaps focuses on the cases where the structure $\mathbf{B}$ is a graph; we refer the reader to the survey [3] for further information and pointers, and also can reference the related articles [13, 10, 11]. The introduction in the survey [3] suggests that the problems $\mathsf{SCSP}(\mathbf{B})$ "seem to be very difficult to classify in terms of complexity", and that "standard methods to prove easiness or hardness fail." Indeed, in contrast to the vanilla CSP, there is no known way to reduce the complexity classification of the problems $\mathsf{SCSP}(\mathbf{B})$ to a classification of clones. In particular, there is no known result showing that the complexity of a problem $\mathsf{SCSP}(\mathbf{B})$ depends only on the relations that are primitive positive definable from $\mathbf{B}$. Thus far, there has been no success in using algebraic information based on the polymorphisms of $\mathbf{B}$ to deduce complexity hardness consequences for the problem $\mathsf{SCSP}(\mathbf{B})$. (The claims given here are relative to the best of our knowledge).

In this article, we give (to our knowledge) the first result which allows one to use algebraic information from the polymorphisms of a structure $\mathbf{B}$ to infer information about the complexity hardness of $\mathsf{SCSP}(\mathbf{B})$. Let us assume that the structures under discussion are finite relational structures. It is known and straightforward to verify that the problem $\mathsf{SCSP}(\mathbf{B})$ polynomial-time reduces to the problem $\mathsf{CSP}(\mathbf{B}^+)$, where $\mathbf{B}^+$ denotes the expansion of $\mathbf{B}$ by constants [3, Section 2]. We give a sufficient condition for the problem $\mathsf{CSP}(\mathbf{B}^+)$ to polynomial-time reduce to the problem $\mathsf{SCSP}(\mathbf{B})$, and hence for the equivalence of these two problems (up to polynomial-time reducibility). From a high level, our sufficient condition requires a certain relationship between the diagonal and the image of an operation, for each operation in the polymorphism clone of $\mathbf{B}$. Any structure $\mathbf{B}$ whose polymorphisms are all essentially unary satisfies our sufficient condition, and a corollary of our main theorem is that, for any such structure $\mathbf{B}$ (having a non-trivial universe), the problem $\mathsf{SCSP}(\mathbf{B})$ is NP-complete. In the classification of two-element structures [7, Proposition 6.11], each structure on which $\mathsf{SCSP}(\mathbf{B})$ is proved NP-complete has only essentially unary polymorphisms (this can be inferred from existing results [6, Theorem 5.1]). Hence, the just-named corollary yields a new algebraic proof of the hardness results needed for this classification; we find this proof to be a desirable, concise alternative to the relational argumentation carried out in previously known proofs of this classification [7, Proposition 6.11], [8, Proposition 4.7].

We hope that our result might lead to further interaction between the study of surjective constraint satisfaction and universal algebra, and in particular that the techniques that we present might be used to prove new hardness results or to simplify known hardness proofs.

2 Preliminaries

For a natural number n, we use $\underline{n}$ to denote the set $\{1, \ldots, n\}$. We use $\wp(B)$ to denote the power set of a set B.

2.1 Logic and computational problems

We make basic use of the syntax and semantics of relational first-order logic. A *signature* is a set of *relation symbols*; each relation symbol R has an associated arity (a natural number), denoted by $\mathrm{ar}(R)$. A *structure* $\mathbf{B}$ over signature σ consists of a *universe* B which is a set, and an interpretation $R^{\mathbf{B}} \subseteq B^{\mathrm{ar}(R)}$ for each relation symbol $R \in \sigma$. In this article, we assume that signatures under discussion are finite, and focus on finite structures; a structure is finite if its universe is finite. When $\mathbf{B}$ is a structure over signature σ, we define $\mathbf{B}^+$ to be the expansion of $\mathbf{B}$ "by constants", that is, the expansion which is defined on signature $\sigma \cup \{C_b \mid b \in B\}$, where each C_b has unary arity and is assumed not to be in σ, and where $C_b^{\mathbf{B}^+} = \{b\}$.

By an *atom*, we refer to a formula of the form $R(v_1, \ldots, v_k)$ where R is a relation symbol, $k = \mathrm{ar}(R)$, and the v_i are variables; by a *variable equality*, we refer to a formula of the form $u = v$ where u and v are variables. A *pp-formula* (short for *primitive positive formula*) is a formula built using atoms, variable equalities, conjunction ($\wedge$), and existential quantification ($\exists$). A *quantifier-free pp-formula* is a pp-formula that does not contain existential quantification, that is, a pp-formula that is a conjunction of atoms and variable equalities. A relation $P \subseteq B^m$ is *pp-definable* over a structure $\mathbf{B}$ if there exists a pp-formula $\psi(x_1, \ldots, x_m)$ such that a tuple $(b_1, \ldots, b_m)$ is in P if and only if $\mathbf{B}, b_1, \ldots, b_m \models \psi$; when such a pp-formula exists, it is called a *pp-definition* of P over $\mathbf{B}$.

We now define the computational problems to be studied. For each structure $\mathbf{B}$, define $\mathsf{CSP}(\mathbf{B})$ to be the problem of deciding, given a conjunction ϕ of atoms (over the signature of $\mathbf{B}$), whether or not there is a map f to B defined on the variables of ϕ such that $\mathbf{B}, f \models \phi$. For each structure $\mathbf{B}$, define $\mathsf{SCSP}(\mathbf{B})$ to be the problem of deciding, given a pair (U, ϕ) where U is a set of variables and ϕ is a conjunction of atoms (over the signature of $\mathbf{B}$) with variables from U, whether or not there is a surjective map $f : U \to B$ such that $\mathbf{B}, f \models \phi$.

Note that these two problems are sometimes formulated as relational homomorphism problems; for example, one can define $\mathsf{SCSP}(\mathbf{B})$ as the problem of deciding, given a structure $\mathbf{A}$ over the signature of $\mathbf{B}$, whether or not there is a surjective homomorphism from $\mathbf{A}$ to $\mathbf{B}$. This is an equivalent formulation: an instance (U, ϕ) of $\mathsf{SCSP}(\mathbf{B})$ can be translated naturally to the structure $\mathbf{A}$ with universe U and where $(u_1, \ldots, u_k) \in R^{\mathbf{A}}$ if and only if $R(u_1, \ldots, u_k)$ is present in ϕ; this structure $\mathbf{A}$ admits a surjective homomorphism to $\mathbf{B}$ if and only if (U, ϕ) is a *yes* instance of $\mathsf{SCSP}(\mathbf{B})$ as we have defined it. One can also naturally invert this passage, to translate from the homomorphism formulation to ours. Let us remark that in our formulation of $\mathsf{SCSP}(\mathbf{B})$, when (U, ϕ) is an instance, it is permitted that U contain variables that are not present in ϕ; indeed, whether or not the instance is a *yes* instance may be sensitive to the exact number of such variables, and this is why this variable set is given explicitly.

We now make a simple observation which essentially says that one could alternatively define $\mathsf{SCSP}(\mathbf{B})$ by allowing the formula ϕ to be a quantifier-free pp-formula, as variable equalities may be efficiently eliminated in a way that preserves the existence of a surjective satisfying assignment.

Proposition 2.1 *There exists a polynomial-time algorithm that, given a pair (W, ϕ) where ϕ is a quantifier-free pp-formula with variables from W, outputs a pair (W', ϕ') where ϕ' is a conjunction of atoms with variables from W' and having the following property: for any structure $\mathbf{B}$ (whose signature contains the relation symbols present in ϕ), there exists a surjective map $f : W \to B$ such that $\mathbf{B}, f \models \phi$ if and only if there exists a surjective map $f' : W' \to B$ such that $\mathbf{B}, f' \models \phi'$.*

Proof. The algorithm repeatedly eliminates variable equalities one at a time, until no more exist. Precisely, given a pair (W, ϕ), it iterates the following two steps as long as ϕ contains a variable equality. The first step is to simply obtain ϕ' by removing from ϕ all variable equalities $u = u$ that equate the same variable,

and then replace (W, ϕ) by (W, ϕ'). The second step is to check if ϕ contains a variable equality $u = v$ between two different variables; if so, the algorithm picks such an equality $u = v$, obtains ϕ' by replacing all instances of v with u, and then replaces (W, ϕ) by $(W \setminus \{v\}, \phi')$. The output of the algorithm is the final value of (W, ϕ). It is straightforwardly verified that this final value has the desired property (by checking that each of the two steps preserve the property). □

2.2 Algebra

All operations under consideration are assumed to be of finite arity greater than or equal to 1. We use $\mathsf{image}(f)$ to denote the image of an operation f. The *diagonal* of an operation $f : B^k \to B$, denoted by $\hat{f}$, is the unary operation defined by $\hat{f}(b) = f(b, \ldots, b)$. Although not the usual definition, it is correct to say that an operation $f : B^k \to B$ is *essentially unary* if and only if there exists $i \in \underline{k}$ such that $f(b_1, \ldots, b_k) = \hat{f}(b_i)$.

When $t^1, \ldots, t^k$ are tuples on B having the same arity m and $f : B^k \to B$ is an operation, the tuple $f(t^1, \ldots, t^k)$ is the arity m tuple obtained by applying f coordinatewise. The entries of a tuple t of arity m are denoted by $t = (t_1, \ldots, t_m)$. Let $P \subseteq B^m$ be a relation, and let $f : B^k \to B$ be an operation; we say that *f is a polymorphism of P* or that *P is preserved by f* if for any choice of k tuples $t^1, \ldots, t^k \in P$, it holds that $f(t^1, \ldots, t^k) \in P$. An operation $f : B^k \to B$ is a *polymorphism* of a structure $\mathbf{B}$ if f is a polymorphism of each relation of $\mathbf{B}$; we use $\mathsf{Pol}(\mathbf{B})$ to denote the set of all polymorphisms of $\mathbf{B}$. It is known that, for any structure $\mathbf{B}$, the set $\mathsf{Pol}(\mathbf{B})$ is a *clone*, which is a set of operations that contains all projections and is closed under composition.

We will make use of the following characterization of pp-definability relative to a structure $\mathbf{B}$.

Theorem 2.2 *[9, 4] A non-empty relation $P \subseteq B^m$ is pp-definable over a finite structure $\mathbf{B}$ if and only if each operation $f \in \mathsf{Pol}(\mathbf{B})$ is a polymorphism of P.*

3 Hardness result

Throughout this section, B will be a finite set; we set $n = |B|$ and use $b_1^*, \ldots, b_n^*$ to denote a fixed enumeration of the elements of B.

We give a complexity hardness result on $\mathsf{SCSP}(\mathbf{B})$ under the assumption that the polymorphism clone of $\mathbf{B}$ satisfies a particular property, which we now define. We say that a clone C on a set B is *diagonal-cautious* if there exists a map $G : B^n \to \wp(B)$ such that:

- for each operation $f \in C$, it holds that $\mathsf{image}(f) \subseteq G(\hat{f}(b_1^*), \ldots, \hat{f}(b_n^*))$, and
- for each tuple $(b_1, \ldots, b_n) \in B^n$, if $\{b_1, \ldots, b_n\} \neq B$, then $G(b_1, \ldots, b_n) \neq B$.

Roughly speaking, this condition yields that, when the diagonal of an operation $f \in C$ is not surjective, then the image of f is contained in a proper subset of B that is given by G as a function of $\hat{f}$.

Example 3.1 When a clone consists only of essentially unary operations, it is diagonal-cautious via the map $G(b_1, \ldots, b_n) = \{b_1, \ldots, b_n\}$, as for an essentially unary operation f, it holds that $\mathsf{image}(f) \subseteq \{\hat{f}(b_1^*), \ldots, \hat{f}(b_n^*)\} = \mathsf{image}(\hat{f})$. □

Example 3.2 When each operation in a clone has a surjective diagonal, the clone is diagonal-cautious via the map G given in the previous example. □

The following lemma is the key to our hardness result; it provides a *quantifier-free* pp-formula which will be used as a gadget in the hardness proof.

Lemma 3.3 *Suppose that* $\mathbf{B}$ *is a finite structure whose universe* B *has size strictly greater than* 1*, and suppose that* $\mathsf{Pol}(\mathbf{B})$ *is diagonal-cautious via* G*. There exists a quantifier-free pp-formula* $\psi(v_1, \ldots, v_n, x, y_1, \ldots, y_m)$ *such that:*

(1) If it holds that $\mathbf{B}, b_1, \ldots, b_n, c, d_1, \ldots, d_m \models \psi$*, then* $b_1, \ldots, b_n, c, d_1, \ldots, d_m \in G(b_1, \ldots, b_n)$*.*

(2) For each $c \in B$*, it holds that* $\mathbf{B}, b_1^*, \ldots, b_n^*, c \models \exists y_1 \ldots \exists y_m \psi$*.*

(3) If it holds that $\mathbf{B}, b_1, \ldots, b_n \models \exists x \exists y_1 \ldots \exists y_m \psi$*, then there exists a unary polymorphism* u *of* $\mathbf{B}$ *such that* $(u(b_1^*), \ldots, u(b_n^*)) = (b_1, \ldots, b_n)$*.*

Proof. Let

$$\begin{array}{ccc} t^1 & = & (t^1_1, \ldots, t^1_{n^n}) \\ \vdots & & \vdots \\ t^n & = & (t^n_1, \ldots, t^n_{n^n}) \end{array}$$

be tuples from $B^{(n^n)}$ such that the following three conditions hold:

(α) It holds that $\{(t^1_i, \ldots, t^n_i) \mid i \in \underline{n^n}\} = B^n$.

(β) For each $i \in \underline{n}$, it holds that $\{t^1_i, \ldots, t^n_i\} = \{b^*_i\}$.

(γ) It holds that $\{t^1_{n+1}, \ldots, t^n_{n+1}\} = B$.

Visualizing the tuples as rows (as above), condition (α) is equivalent to the assertion that each tuple from B^n occurs exactly once as a column; condition (β) enforces that the first n columns are the tuples with constant values $b^*_1, \ldots, b^*_n$ (respectively); and, condition (γ) enforces that the $(n+1)$th column is a *rainbow* column in that each element of B occurs exactly once in that column.

Let P be the (n^n)-ary relation $\{f(t^1, \ldots, t^n) \mid \ f \text{ is an } n\text{-ary polymorphism of } \mathbf{B}\ \}$. It is well-known and straightforward to verify that the relation P is preserved by all polymorphisms of $\mathbf{B}$. By Theorem 2.2, we have that P has a pp-definition $\phi(w_1, \ldots, w_{n^n})$ over $\mathbf{B}$. We may and do assume that ϕ is in prenex normal form, in particular, we assume $\phi = \exists z_1 \ldots \exists z_q \theta(w_1, \ldots, w_{n^n}, z_1, \ldots, z_q)$ where θ is a conjunction of atoms and equalities.

Since $t^1, \ldots, t^n \in P$, there exist tuples $u^1, \ldots, u^n \in B^q$ such that, for each $k \in \underline{n}$, it holds that $\mathbf{B}, (t^k, u^k) \models \theta$. By condition (α), there exist values $a_1, \ldots, a_q \in \underline{n^n}$ such that, for each $i \in \underline{q}$, it holds that $(u^1_i, \ldots, u^n_i) = (t^1_{a_i}, \ldots, t^n_{a_i})$. Define $\psi(w_1, \ldots, w_{n^n})$ as $\theta(w_1, \ldots, w_{n^n}, w_{a_1}, \ldots, w_{a_q})$. We associate the variable tuples $(w_1, \ldots, w_{n^n})$ and $(v_1, \ldots, v_n, x, y_1, \ldots, y_m)$, so that ψ may be viewed as a formula with variables from $\{v_1, \ldots, v_n, x, y_1, \ldots, y_m\}$. We verify that ψ has the three conditions given in the lemma statement, as follows.

(1): Suppose that $\mathbf{B}, b_1, \ldots, b_n, c, d_1, \ldots, d_m \models \psi$. Then $(b_1, \ldots, b_n, c, d_1, \ldots, d_m)$ is of the form $f(t^1, \ldots, t^n)$ where f is a polymorphism of $\mathbf{B}$. We have

$$\{b_1, \ldots, b_n, c, d_1, \ldots, d_m\} \subseteq \mathsf{image}(f) \subseteq G(\hat{f}(b^*_1), \ldots, \hat{f}(b^*_n)) = G(b_1, \ldots, b_n).$$

The second containment follows from the definition of *diagonal-cautious*, and the equality follows from (β).

(2): We had that, for each $k \in \underline{n}$, it holds that $\mathbf{B}, (t^k, u^k) \models \theta$. By the choice of the a_i and the definition of ψ, it holds (for each $k \in \underline{n}$) that $\mathbf{B}, t^k \models \psi$. Condition (2) then follows immediately from conditions (α) and (β).

(3): Suppose that $\mathbf{B}, b_1, \ldots, b_n \models \exists x \exists y_1 \ldots \exists y_m \psi$. By definition of ψ, we have that there exists a tuple beginning with $(b_1, \ldots, b_n)$ that satisfies θ on $\mathbf{B}$. By the definition of θ, we have that there exists a tuple t beginning with $(b_1, \ldots, b_n)$ such that $t \in P$. There exists a polymorphism f of $\mathbf{B}$ such that $t = f(t^1, \ldots, t^n)$. By condition (β), we have that $(\hat{f}(b_1^*), \ldots, \hat{f}(b_n^*)) = (b_1, \ldots, b_n)$. $\square$

Let us make some remarks. The relation P in the just-given proof is straightforwardly verified (via Theorem 2.2) to be the smallest pp-definable relation (over $\mathbf{B}$) that contains all of the tuples $t^1, \ldots, t^n$. The definition of ψ yields that the relation defined by ψ (over $\mathbf{B}$) is a subset of P; the verification of condition (2) yields that each of the tuples $t^1, \ldots, t^n$ is contained in the relation defined by ψ. Therefore, the formula ψ defines precisely the relation P. A key feature of the lemma, which is critical for our application to surjective constraint satisfaction, is that the formula ψ is quantifier-free. We believe that it may be of interest to search for further applications of this lemma.

The following is our main theorem.

Theorem 3.4 *Suppose that* $\mathbf{B}$ *is a finite structure such that* $\mathsf{Pol}(\mathbf{B})$ *is diagonal-cautious. Then the problem* $\mathsf{CSP}(\mathbf{B}^+)$ *many-one polynomial-time reduces to* $\mathsf{SCSP}(\mathbf{B})$.

Proof. The result is clear if the universe B of $\mathbf{B}$ has size 1, so assume that it has size strictly greater than 1. Let $\psi(v_1, \ldots, v_n, x, y_1, \ldots, y_m)$ be the quantifier-free pp-formula given by Lemma 3.3. Let ϕ be an instance of $\mathsf{CSP}(\mathbf{B}^+)$ which uses variables U. The reduction creates an instance of $\mathsf{SCSP}(\mathbf{B})$ as follows. It first creates a quantifier-free pp-formula ϕ' that uses variables

$$U' = U \cup \{v_1, \ldots, v_n\} \cup \bigcup_{u \in U} \{y_1^u, \ldots, y_m^u\}.$$

Here, each of the variables given in the description of U' is assumed to be distinct from the others, so that $|U'| = |U| + n + |U|m$. Let $\phi^=$ be the formula obtained from ϕ by replacing each atom of the form $C_{b_j^*}(u)$ by the variable equality $u = v_j$. The formula ϕ' is defined as $\phi^= \wedge \bigwedge_{u \in U} \psi(v_1, \ldots, v_n, u, y_1^u, \ldots, y_m^u)$. The output of the reduction is the algorithm of Proposition 2.1 applied to (U', ϕ').

To prove the correctness of this reduction, we need to show that there exists a map $f : U \to B$ such that $\mathbf{B}^+, f \models \phi$ if and only if there exists a surjective map $f' : U' \to B$ such that $\mathbf{B}, f' \models \phi'$.

For the forward direction, define $f^= : U \cup \{v_1, \ldots, v_n\} \to B$ to be the extension of f such that $f^=(v_i) = b_i^*$ for each $i \in \underline{n}$. It holds that $f^=$ is surjective and that $\mathbf{B}, f^= \models \phi^=$. By property (2) in the statement of Lemma 3.3, there exists an extension $f' : U' \to B$ of $f^=$ such that $\mathbf{B}, f' \models \phi'$.

For the backward direction, we argue as follows. We claim that $\{f'(v_1), \ldots, f'(v_n)\} = B$. If not, then by the definition of diagonal-cautious, it holds that $G(f'(v_1), \ldots, f'(v_n)) \neq B$; by property (1) in the statement of Lemma 3.3 and by the definition of ϕ', it follows that $f'(u') \in G(f'(v_1), \ldots, f'(v_n))$ for each $u' \in U'$, contradicting that f' is surjective. By property (3) in the statement of Lemma 3.3, there exists a unary polymorphism u of $\mathbf{B}$ such that $(u(b_1^*), \ldots, u(b_n^*)) = (f'(v_1), \ldots, f'(v_n))$; by the just-established claim, u is a bijection. Since the set of unary polymorphisms of a structure is closed under composition and since B is by assumption finite, the inverse u^{-1} of u is also a polymorphism of $\mathbf{B}$. Hence it holds that $\mathbf{B}, u^{-1}(f') \models \phi'$, where $u^{-1}(f')$ denotes the composition of f' with u^{-1}. Since $u^{-1}(f')$ maps each variable v_j to b_j^*, we can infer that $\mathbf{B}^+, u^{-1}(f') \models \phi$. $\square$

Corollary 3.5 *Suppose that* $\mathbf{B}$ *is a finite structure whose universe* B *has size strictly greater than* 1*. If each polymorphism of* $\mathbf{B}$ *is essentially unary, then* $\mathsf{SCSP}(\mathbf{B})$ *is NP-complete.*

Proof. The problem $\mathsf{SCSP}(\mathbf{B})$ is in NP whenever $\mathbf{B}$ is a finite structure, so it suffices to prove NP-hardness. By Example 3.1, we have that $\mathsf{Pol}(\mathbf{B})$ is diagonal-cautious. Hence, we can apply Theorem 3.4, and it suffices to argue that $\mathsf{CSP}(\mathbf{B}^+)$ is NP-hard. Since $\mathbf{B}^+$ is by definition the expansion of $\mathbf{B}$ with constants, the polymorphisms of $\mathbf{B}^+$ are exactly the idempotent polymorphisms of $\mathbf{B}$; here then, the polymorphisms of $\mathbf{B}^+$ are the projections. It is well-known that a structure having only projections as polymorphisms has a NP-hard CSP [5] (note that in this case, Theorem 2.2 yields that every relation over the structure's universe is pp-definable). □

Acknowledgements. The author thanks Matt Valeriote, Barny Martin, and Yuichi Yoshida for useful comments and feedback. The author was supported by the Spanish Project FORMALISM (TIN2007-66523), by the Basque Government Project S-PE12UN050(SAI12/219), and by the University of the Basque Country under grant UFI11/45.

References

[1] L. Barto and M. Kozik. Constraint satisfaction problems of bounded width. In *Proceedings of FOCS'09*, 2009.

[2] Manuel Bodirsky. Complexity classification in infinite-domain constraint satisfaction. *CoRR*, abs/1201.0856, 2012.

[3] Manuel Bodirsky, Jan Kára, and Barnaby Martin. The complexity of surjective homomorphism problems - a survey. *Discrete Applied Mathematics*, 160(12):1680–1690, 2012.

[4] V. G. Bodnarčuk, L. A. Kalužnin, V. N. Kotov, and B. A. Romov. Galois theory for post algebras, part I and II. *Cybernetics*, 5:243–252, 531–539, 1969.

[5] A. Bulatov, P. Jeavons, and A. Krokhin. Classifying the Complexity of Constraints using Finite Algebras. *SIAM Journal on Computing*, 34(3):720–742, 2005.

[6] Hubie Chen. A rendezvous of logic, complexity, and algebra. *ACM Computing Surveys*, 42(1), 2009.

[7] N. Creignou, S. Khanna, and M. Sudan. *Complexity Classification of Boolean Constraint Satisfaction Problems*. SIAM Monographs on Discrete Mathematics and Applications. Society for Industrial and Applied Mathematics, 2001.

[8] Nadia Creignou and Jean-Jacques Hébrard. On generating all solutions of generalized satisfiability problems. *ITA*, 31(6):499–511, 1997.

[9] D. Geiger. Closed Systems of Functions and Predicates. *Pacific Journal of Mathematics*, 27:95–100, 1968.

[10] Petr A. Golovach, Daniël Paulusma, and Jian Song. Computing vertex-surjective homomorphisms to partially reflexive trees. *Theor. Comput. Sci.*, 457:86–100, 2012.

[11] Pavol Hell. Graph partitions with prescribed patterns. *Eur. J. Comb.*, 35:335–353, 2014.

[12] P. Idziak, P. Markovic, R. McKenzie, M. Valeriote, and R. Willard. Tractability and learnability arising from algebras with few subpowers. *SIAM J. Comput.*, 39(7):3023–3037, 2010.

[13] Hannes Uppman. Max-sur-csp on two elements. In *CP*, pages 38–54, 2012.